Acceptable Cancer Risks: Probabilities and Beyond

Paolo F. Ricci
University of California, Los Angeles
Los Angeles, California

Louis Anthony Cox, Jr.
U.S. West Advanced Technologies
Englewood, Colorado

John P. Dwyer
University of California
Berkeley, California

The acceptability of cancer risk requires consideration of factors that extend beyond mere numerical representations, such as either individual lifetime risk in excess of background and excess incidence. Recently, use of these numbers has been tempered by the addition of qualitative weights-of-evidence that describe the degree of support provided by animal and epidemiologic results. Nevertheless, many other factors, most of which are not quantitative, require incorporation but remain neglected by the analyst eager to use quantitative results.

In this paper we show that simple risk measures are often fraught with problems. Moreover, these measures do not incorporate the very essence of acceptability, which includes notions of responsibility, accountability, equity, and procedural legitimacy, among others. We link the process of risk assessment to those legal and regulatory standards that shape it. These standards are among the principal means to resolve risk-related disputes and to enhance the balancing of competing interests when science and law meet on uncertain and often conjectural ground.

We conclude the paper with a proposal for the portfolio approach to manage cancer risks and to deal with uncertain scientific information. This approach leads to the concept of "provisional acceptability," which reflects the choices available to the decisionmaker, and the trade-offs inherent to such choices.

Agencies, industry, and the public demand clear standards for judging the acceptability of risks. Numerical values could reduce debate and ambiguity, clarify the responsibilities of businesses, and provide data for regulatory, judicial, and legislative deliberations.[1,2]

Recognizing that a single risk level is not appropriate in all contexts, it is tempting to propose specific numerical "acceptability" values for different classes of risks.[3] For example, the average acceptable excess individual lifetime fatality probability of cancer from occupational exposures (assuming full disclosure and informed consent) might be set at 1×10^{-4}. A level of 1×10^{-6} could be defined as acceptable for the general public experiencing involuntary exposures. A much higher risk level, such as 1×10^{-3}, might be appropriate for sales of inherently dangerous products to fully informed, willing customers. Aggregate incidence could be acceptable if it were less than some value, for example, unity. However, even such a range of numbers over different contexts is neither conceptually adequate nor sufficient as a basis for responsible decisionmaking.

Whether a risk is "acceptable" generally depends not only on its objective quantitative probability and the nature and severity of the consequences, but also on societal and political factors. Single numerical estimates of individual and population risks do not incorporate those qualitative aspects of risk. Protection of individual rights, the equity of risk-benefit distribution, prudence when facing uncertainty, the absence of knowledge, the legitimacy of the risk management process, and public attitudes toward and perceptions of risks do not lend themselves well to bare numerical representations.[4]

This paper examines these issues, assesses some current approaches to social and legal risk management, and proposes a risk-portfolio approach in which risk acceptability is an evolving concept. We begin with three concepts of risk.

Individual and Population Risks

Two related concepts are useful in describing risk to an individual: the *total risk* to an individual of a particular adverse health consequence, such as cancer; and the concept of an *attributable risk* describing the incremental contribution to total risk made by a particular source or cause (e.g., the contribution made by cigarette smoking to the risk of lung cancer). Finally, we discuss *population risk*, in which individual risk is aggregated over the population at risk.

Total Individual Risk

The total risk to an individual of developing some undesirable health response, such as death from cancer, may be defined as the probability that he will develop the response in a given year t, if he has survived until then. This probability is also called the individual's discrete time "hazard rate"

for the response in year t. Hazard rates can be used to calculate probabilities of cause-specific deaths or illnesses, to derive survival time probability distributions, and to quantify total risks over time.[5,6] Individual hazard rates for chronic health effects typically depend on the exogenous factors to which an individual has been exposed, including the extent of exposure to a particular chemical or radiation. Endogenous factors such as the efficiency of the body's repair mechanisms, genetic predisposition toward response, and so forth, may also affect individual hazard functions. Endogenous factors mediating between exposures and health responses usually vary widely across individuals and often cannot be observed. Thus, individual responses to specific exposures are quite heterogeneous. Even if an individual knew his own exposure history to a chemical, he would generally remain uncertain about his own future hazard function, and hence, about his probability of adverse response.

A problem often overlooked in discussing risk numbers is that interpreting them in terms of expected annual frequencies, or average times until occurrence, can be misleading. For example, a leukemia hazard rate from exposure to benzene in the workplace of one expected excess case per million person-years of exposure does not imply that the probability of a randomly selected worker developing cancer from a year of exposure is 1×10^{-6}. For any individual, it is considerably more likely than not that the actual waiting time to the first arrival will be less than the average (or "expected") waiting time.[5] The probability that a randomly selected individual will develop cancer from a year of exposure is actually 6.3×10^{-5}. The actual individual risk may thus exceed what was believed to be acceptable under a simple regulatory scheme that requires acceptable exposure to be determined from an average risk of, say, 1×10^{-6}.

Attributable Risk

In practical risk assessment and management, the problem is not how to estimate an individual's *total risk* of some health response, but how to estimate the *incremental* contribution to his risk made by some particular cause or source. This is the risk that is said to be *attributable* to the source.[6,7] Few concepts in risk analysis have occasioned as much perplexity and debate as that of the risk attributable to a source.[8,9]

One elementary model that can clarify the meaning of attributability postulates that each source of risk "competes" with other sources to be the first to cause an adverse response. Suppose that N sources contribute to the risk in an individual. Each source can be thought of as firing a random stream of biologically effective molecules that cause "hits" in the exposed individual. The average arrival rate or intensity of hits from source i at time t is the source's hazard rate at time t, denoted by $h_i(t)$. If the N potential sources of a health effect are statistically independent so that the arrival rate of hits from one source is unaffected by the presence of other sources, then the total risk to the individual at time t from all sources is given by the sum $h(t) = h_i(t) + \ldots + h_n(t)$. Since any source i will contribute the fraction $h_i/h(t)$ of all expected hits per unit time at time t, this ratio equals the probability that source i contributed the hit that caused the observed health response, and $h_i(t)$ is the risk attributable to source i at time t.

This "competing risk" definition of attributable risk is satisfactory only when the random arrival model correctly represents the nature of causation. For example, suppose that occurrence of a health effect depends on whether the total number of hits received from all sources within a certain amount of time exceeds a certain threshold. Then, if a response occurs, it cannot even in principle be ascribed to any single source.[10] Similarly, if the presence of factor A doubles the hazard rate from factor B, then a hit from B may be partly blamed on A. In such cases of joint and multiple causation, assignment of shares of causation to the different contributing sources is as much a matter of policy as one of science.[8,10]

Population Risk

Individual risks are not sufficient to determine the effect that a risk management choice can have on those at risk. The full impact of a choice can only be evaluated by looking at the *distribution* of effects in the affected population as a whole.

If the population at risk consists of several "types" of individuals, with each type corresponding to a homogenous subpopulation of individuals having identical hazard functions for death from cancer, then in each homogenous subpopulation the amounts of time (number of life-years) that the members have left until death will be statistically independent, identically distributed random variables. At any time, the total number of remaining life-years in the population, summed over all the individuals now in it, will be approximately normally distributed, with mean and variance equal to the sums of the means and variances, respectively, of the remaining life-years in each subpopulation. The problem of evaluating population risk in large populations can thus be reduced to the problem of evaluating normal distributions for the attribute "remaining life-years in the population." This is a standard decision-analytic problem.[11-13]

If an individual considers himself to be a randomly selected member of the population, then his individual risk is the expected value of the risks for all individuals in the population. There is a paradox: if two different population risk distributions have identical means, then every individual should be indifferent among them (based on his own expected risk). But the distributions may not be equally desirable from a societal perspective. In fact, a variety of different distributions of risks and uncertainties can occur at the population level. It is not the expected number of occurrences per million person-years of exposure in the whole population that counts in determining equity, but the way this risk density is distributed among identifiable subgroups of the population. Consider choosing among the following situations:[10]

Case A (*Uniform population risk*): Each of 100 people independently is exposed to a 0.01 chance of disease.

Case B (*Anonymous sensitive subpopulation*): 10 of the 100 people are exposed to a 0.10 chance of disease. The rest have a zero chance of disease. Now one knows which type he is unless and until he gets the disease.

Case C (*Known high risk population*): 50 of the 100 people (e.g., neighbors living within a certain distance of an industrial facility) are at high risk and know it. Each of these individuals has a 0.02 probability of disease; the remaining 50 people have zero risk.

Case D (*Uncertain individual risk*): Each individual has a random probability, independently drawn from a uniform distribution between 0 and 0.02, of getting the disease.

Case E (*Uncertain population risk*): Each individual independently has the same probability of getting the disease. The magnitude of this probability is uncertain; however, it is judged equally likely to be anywhere between 0 and 0.02.

Imagine trying to rank these situations in terms of relative social desirability. Simply using the expected number of casualties as a summary of population risk results in the same assessment of risk (one expected case) for all of these examples. Aggregating risk obscures uncertainties, heterogeneities, and inequities. The five cases involve important trade-offs: between number of people exposed and magni-

tude of exposure per person, between certain and uncertain risks, and between equitably and inequitably distributed risks.

The concepts of individual and aggregate risks are quantitative expressions that, at best, summarize all relevant biological information. However important these concepts are, there are other social and ethical dimensions that also contribute to risk acceptability. For example, what constitutes an acceptable risk for one person to offer to another depends on both rights (e.g., the right of free choice to assume an otherwise "unacceptable" risk) and duties (e.g., the legal duty not to endanger others even with their consent).

Acceptability also depends on the responsibilities that acceptance entails. For example, an informed cancer patient, cognizant of the risks and uncertainties associated with a new, unlicensed chemotherapeutic drug, might consider its administration personally acceptable but not acceptable to administer to someone else. At the other extreme, the social decisionmaker faces accountability: if she regulates the specific risk by selecting a specific option, but other unknown risks eventually result from her choice, even though she may not have been *responsible* for those risks, she may be held socially *accountable* for them. The next section explores these concepts.

Attitudes Toward Risk

The issue of voluntary acceptance underlies many debates about acceptability in the context of making decisions about risks. In a society that values individual choice, a risk that an individual is willing to take for himself may be acceptable, even though a quantitatively similar risk imposed by another is not.

The concept of voluntary acceptance of risks has several important components. First, there are gradations of "volluntariness." A truly voluntary activity is one that an individual is free and able to reject without penalty, that he can control while undertaking, and whose risks are fully known or easily discoverable. Less clearly voluntary are activities that an individual is intially free to reject without penalty, but over which he loses some control over time. Drinking, smoking, and other addictive activities are examples. Voluntariness may be further compromised when advertisements and deliberate inducements play a major role in the initial decision to participate in the activity.[14]

The appearance of controllability also plays a large role in individual judgments of risk acceptability. So does the belief that a truly free choice requires participants to have full information about what is being chosen. Right-to-know, need-to-know, and duty-to-warn legislation, as well as principles governing "adhesion" contracts, make clear the notion that lack of information about the risks in an economic transaction undermines one's consent to accept risk.

Individual perceptions and preferences regarding risks can be fragile, changeable, overly sensitive to initial impressions, and unreliable. This raises crucial questions for policymakers. When public perceptions and statistical realities conflict, is it the decisionmaker's duty to represent the views and preferences of members of society, or to protect what he considers to be in their true interests? How can a line be drawn between responsibility and paternalism when public preferences appear to be based on inaccurate perceptions? Does the judgment that a population risk is acceptable apply to the median, the average, the most highly susceptible, or the most risk-averse individual in the population?

In practice, the distribution of risks and risk perceptions within the population at risk is often left unexplored. Instead, the individual risk to the maximally exposed individual is given special attention by regulatory agencies. The assumption is that if risk to the maximally exposed individual is acceptably small, then the risk to the entire population is plausibly also acceptable. However, as discussed above, this view is suspect: population risk and individual risk should be assessed separately. For example, a maximally exposed adult may have a hazard rate that is higher than a child's, even though both are identically exposed to a particular toxic chemical.

In addition, the alternative concept of a "maximally threatened" individual is poorly defined. For example, who is more "threatened," an individual whose hazard rate is increased from 0.01 to 0.02, or an individual whose hazard rate is increased from 0.02 to 0.03? The former faces 50 years of lost life expectancy (with it being considerably more likely than not that the loss will be greater than its expected value), while the latter faces only a 17 year reduction in life expectancy. On the other hand, with less expected life left, the second individual might be thought to suffer more for each additional expected year lost.

Procedural Legitimacy

Risk decisions usually involve numerous parties with conflicting interests. Individuals may be considered to have accepted a risk "voluntarily" if they "agree to" the decisionmaking process leading to it. Very often the perceived legitimacy of that process depends on the degree to which the risk-bearing public could participate in regulatory decisions. The history of nuclear power and hazardous facility siting in this country illustrates the close connection between perceptions of voluntary participation (or opportunity to participate) in judging risk acceptability, the perceived legitimacy of regulatory processes, and the acceptability of risks.[14,15]

An individual may agree to abide by the results of a social decision process because she expects to gain from the process on average, even though she may lose from particular decisions. From this perspective, acceptability of each risky prospect is not the most relevant issue. Rather, the acceptability and equity of the entire *portfolio* of risks, selected through the process over time, are what matter.[16] We will develop this theme by examining the regulatory and judicial process by which our society's risk activity portfolio is largely determined in practice, with focus on its fairness when methods and results are at the frontiers of knowledge and when societal risks, economic costs, and benefits may be large.

Private Litigation

Private litigation over environmental injuries have been standard fare for hundreds of years. William Aldred, to take one well known example, successfully sued his neighbor for damages on the ground that the neighbor's pigsty "corrupted" the air and thereby prevented Aldred from living in his own home.[17] From these humble common law beginnings, courts and legislatures have developed a massive body of private law to compensate environmental injuries by awarding damages and to prevent future injuries through both injunctive relief and the deterrent effect of potential damage awards.

Although private litigation is no longer the principal means to regulate environmental injuries, it continues to play a potentially important role, especially with regard to its principal goals of compensation and deterrence. Private litigation is almost the sole means, other than first party insurance, to obtain monetary compensation for personal injuries and property damage, something that most regulatory schemes fail to do (although they could). In addition, many regulatory schemes are not comprehensive in scope; the legislature and the agency have failed to regulate (or failed to regulate adequately) significant environmental problems. In these cases, private litigation could supplement or fill in the gaps in existing regulatory programs.

Finally, nuisance and other tort doctrines often allow plaintiffs to secure prompt injunctive relief, such as a temporary restraining order, to prevent imminent injuries. A judge will grant relief after making an ad hoc balancing of the equities in the particular case. Under most regulatory schemes, by contrast, an injured or threatened citizen must first persuade a government official to seek injunctive relief from the court.

Despite the venerable origins of private environmental litigation, and its unique strengths, there are serious questions about the appropriateness of private remedies for polycentric social questions, and the capacity of courts to decide environmental tort cases. Although some of the following problems could be cured by legislative reforms,[18–20] others are more resistant.

Statutes of limitations. Statutes of limitations traditionally have posed an impassable obstacle to recovery if the injury does not manifest itself for many years. Many states, however, have reduced this obstacle by postponing the time limits for filing suit until the injury manifests itself or until the link between exposure and injury would be known to the "reasonable person," even though it may not be known to the actual plaintiff.

Litigation costs. Private litigation is expensive. Many lawyers are unwilling to undertake a private suit with a low and uncertain probability of success, even though the suit may have merit. The result is substantial undercompensation and reduced deterrent effect.

Complexity. Environmental problems often involve complex, uncertain, and sometimes unresolvable scientific issues. Is it realistic to expect that judges and juries will be able to intelligently address and resolve the technical and scientific questions when science cannot provide defensible answers? How will firms be able to make informed, long-term investments if the outcomes of cases are so uncertain?

Cases alleging environmental injury may take several years to resolve. In part, this is a result of the scientific complexity of the assumptions, theories, and data underlying the disputed factual issues. In part, lengthy delays occur because the judicial system does not have the administrative capacity to handle mass-injury cases, such as injuries from exposure to asbestos. More fundamentally, there has been a basic shift in the focus of tort law. Environmental tort cases today are not bipolar disputes, such as the one between William Aldred and his neighbor. Instead they affect:

> great aggregations of people and vast economic and social interests. The decisions in these cases are preoccupied ... with advancing public control of large-scale activities and altering both the distribution of power and the nature of social values. In such cases, the parties are ... mere placeholders for these larger social interests.[21]

What should a court do if the requested injunctive relief would injure numerous persons who depend on the polluter for employment, taxes, and consumer products? Since these cases raise policy issues involving the allocation of resources among different segments of the community, should those issues be decided in court or left to the political process?

Causation. Causation may be an intractable issue in toxic tort law.[22–25] There are really two causation problems, known as the *indeterminate defendant* and the *indeterminate plaintiff*.[21] The indeterminate defendant problem occurs because in most cases there are numerous sources of pollution, which makes it difficult to directly identify the source (or sources) that caused the plaintiff's injury. Some courts have addressed this problem by apportioning liability among defendants through a market share theory, where the harm results from an identifiable product.[26] This approach has no applicability where the injury may have resulted from a variety of ubiquitous chemicals dispersed through the environment, some of which act synergistically. The market share theory also may be of little help when, as is common in toxic torts, the injury manifests itself years after the injurious exposure occurred, and the potential plaintiffs have gone out of business or are difficult to identify.

The indeterminate plaintiff problem arises because many injuries (particularly injuries to health) are "nonsignature," meaning that the injuries could have resulted from a number of causes, some of which are natural. For example, although several members of the community may have evidence that the chemicals in the defendant's air emissions cause lung cancer, comparisons with other communities suggest that most of the lung cancer cases would have developed regardless of the defendant's polluting activity. In the *Agent Orange* case, which began in 1975 and concluded in 1984, the trial judge granted summary judgment to several chemical companies against non-settling plaintiffs, largely on the trial judge's conclusion that the plaintiffs had failed to establish this kind of causation. Although the court also approved a $180 million settlement covering most plaintiffs, the court's approval was based more on sympathy for the veterans than on belief that they had demonstrated causation.[27,28] The problem is easy to state but difficult to resolve:

> *identification*, ordinarily a routine issue of cause in fact at common law, is a costly enterprise that relies on types of evidence and probability judgments which can be regarded as ill-suited to traditional resolution through the adversary process.[20]

One problematic aspect of the indeterminate plaintiff issue is illustrated by the following example. Suppose there was a clear correlation between the defendant's polluting activity and a 15 percent increase in the expected number of lung cancer cases in the vicinity of the defendant's plant (of course, such correlations are almost never clear cut for nonsignature diseases), but that it was unclear which cases resulted from the defendant's air emissions. Should damages be awarded in full to every person with lung cancer? Or is the proper award to each person with lung cancer 1/115 of a sum of money representing damages for 15 cases (i.e., 15 full awards spread out over every 115 people with lung cancer)? It is doubtful in these circumstances that it is ever possible to avoid both overcompensation and undercompensation.

Agency Decisionmaking

In principle, a regulatory system should overcome many of the shortcomings of private litigation. Issues of expertise, complexity, political accountability, and causation should prove to be smaller obstacles. Nevertheless, pervasive scientific uncertainty, as well as the ambiguity inherent in statutory commands, make agency decisionmaking problematic as well.

A regulatory agency generally uses one of two processes to manage risks: administrative adjudication or rulemaking. Administrative adjudication is the case-by-case determination of regulatory issues, such as whether to cancel the registration of a particular pesticide. In adjudication, an administrative law judge presides over a formal, trial-type hearing and, after listening to evidence, formulates findings of law and fact. Rulemaking, by contrast, produces regulations governing the activities of entire industries. Most federal rulemaking proceedings under environmental and worker safety statutes do not employ live hearings, but instead rely on written submissions that constitute the administrative record. Generally, most agency policies and decisions concerning health and safety are generated through rulemaking proceedings.

One of the hallmarks of modern agency decisionmaking in the United States is extensive public participation.[29] Before adopting a regulation, for example, the agency must publish

the proposed regulation, a summary of the reasons and factual bases for the regulation, and the time and place for submitting written comments. Anyone may submit written comments, reports, data, or objections related to the agency's proposed rule, and any documents submitted to the agency are available for public review. In publishing a final regulation, the agency must explain the basis for the rule, including any changes from the proposed rule, and must respond to (but need not rely on) all material comments and objections.

Public participation serves multiple, sometimes conflicting purposes. At one level, public participation legitimates agency decisionmaking. Many agency decisions are controversial because there are not adequate data leading unambiguously to a single policy choice. Such decisions are necessarily value-laden. Public participation, coupled with the requirement that the agency take seriously and respond to all material comments, can ensure that all interested parties have roughly equal access to political decisionmaking, and thus can help to ensure at least a minimal level of public accountability by agency officials.

At another level, unrestricted public participation may also help to ensure that the agency has made the "best" decision. By being required to consider all points of view, the agency is much less likely to overlook relevant data or perspectives, and also is more likely to be able to overcome its institutional biases.[30]

In addition to procedural constraints on agency decisionmaking, the legislature also sets the substantive criteria for agency regulations in virtually all risk management statutes. One characteristic of modern environmental and worker safety statutes is their relative specificity.

The two extremes of pollution control standards are "health-based" and "technology-based" standards. An example of a "health-based" standard is section 112 of the Clean Air Act. Under that provision, the Administrator of the Environmental Protection Agency must set a hazardous air pollutant standard that "in his judgment provides an ample margin of safety to protect the public health." The current judicial interpretation of this language prohibits the Administrator from considering implementation costs or technological feasibility in setting an "acceptable" or "safe" level of risk.[31] After fixing this "safe" risk level, the agency must further reduce the emission rate (to add the statutory ample margin of safety) to the extent permitted by implementation costs and technological feasibility. Thus, under this interpretation, even the best available control technologies do not necessarily meet the statutory standard.

Not surprisingly, this interpretation has caused consternation in regulatory circles, for it requires the agency to determine a level of acceptable risk wholly out of context; implementation costs and technological feasibility ostensibly are irrelevant to the agency's initial determination of acceptable risk. There is some reason to believe that the EPA will seek to circumvent the statutory restriction. The agency recently proposed to define acceptable risk as equivalent to a maximum individual lifetime risk in the neighborhood of 1×10^{-4}. The precise number would vary depending on unspecified factors; in one example, the risk level was 6×10^{-3}.[32] By proposing a relatively high level of acceptable risk—one that few people would define as acceptable—EPA effectively would be able to give much greater weight to costs and feasibility considerations than if it had begun with a more conventional estimate of acceptable risk.

An example of a "technology-based" standard is section 111 of the Clean Air Act. That provision requires the EPA Administrator to set emission standards for new sources of nonhazardous air pollutants that "reflect the degree of emission limitation . . . achievable through application of the best technological system . . . which (taking into consideration the cost of achieving such emission reduction . . .) the Administrator determines has been adequately demonstrated." Thus, in contrast to section 112, section 111 requires EPA to consider implementation costs and the availability of technological controls in setting emission standards.

As these two examples illustrate, the substantive statutory language often is imprecise, and thus leaves the agency considerable room for interpretation. How expansive is the definition of "public health"? Should the standards be set to protect every vulnerable individual, no matter how hypersensitive? If not, where should the line be drawn? How much of a margin is an "ample margin of safety"? How should the Administrator "consider" costs under section 111? How should the Administrator determine what is the "best" system of pollution control? When is a technological control system "adequately demonstrated"? Plainly, in these provisions the legislature has only defined the broad outlines of regulatory policy. Crucial substantive details remain to be worked out in agency rulemaking proceedings.

Despite the inherent ambiguity of the statutory language, however, it is evident that the health-based and technology-based criteria represent two very different approaches to the regulation of human health risks. Even though the statutory language leaves the agency considerable policymaking discretion, these criteria impose important limits on the ambit of agency authority. Health-based criteria require the regulator to focus exclusively on health risks—e.g., their nature, magnitude, and distribution—in determining acceptable emissions standards. Technology-based criteria, by contrast, require the regulator to focus entirely on the availability and cost of pollution control technology. Thus, the legislature, through the substantive statutory criteria for pollution standards, defines the factors that contribute to a determination of acceptable risk in different circumstances.

Judicial Review

A critical element of the rulemaking process is judicial review. Roughly speaking, judicial review is designed to ensure that the agency has conformed to the statute's procedural and substantive requirements. If the agency failed to comply with the statutory requirements, the court would remand the case to the agency for reconsideration in compliance with the statute. The policy goal of widespread public participation in agency decisionmaking is advanced here by statutes and doctrines permitting any person to seek judicial review of agency actions.

Ensuring that the agency has conformed to the procedural requirements is essential to any notion of public participation and agency accountability. Without adequate judicial review, agency officials would be free to disregard conflicting perspectives and to make policy choices without a thorough public airing. For example, in the course of deciding not to regulate formaldehyde under the Toxic Substances Control Act, EPA officials held private meetings with representatives of industry.[33] As a result, opposing views were largely left out of the decisionmaking process, EPA deviated arbitrarily from the existing cancer policy guidelines, and the agency did not follow procedures for internal review of policy decisions. Judicial review of agency procedures could have helped ensure the legitimacy and integrity of the quasi-political process that is central to agency rulemaking.

More commonly, judicial review focuses on whether the agency's substantive decision conforms with the statutory criteria. Inevitably, this issue is linked with questions involving the extent to which the reviewing court should defer to the agency's interpretation of the statute, and the degree to which the court should probe the validity of the agency's technical judgments.

The courts' willingness to read statutes instrumentally, and thus to disregard agency interpretations, was illustrated in *Industrial Union Dep't, AFL-CIO v. American Petroleum Institute*.[34] In that case, a plurality of the Supreme Court remanded the OSHA benzene exposure standard of 1

ppm to the agency for reconsideration. The Court's decision was based on its reading of a convoluted statute as requiring the agency to show that the health risks to workers were "significant," a word that does not appear in the statute. It is plain from the Court's opinion that its interpretation was driven by concerns that OSHA's interpretation would result in excessively strict regulations.[35]

More recently, the federal courts have assumed a deferential posture toward agency interpretation of regulatory statutes. In a series of cases over the last five years, the Supreme Court has emphasized that for reasons of political accountability and technical expertise, regulatory agencies have the primary responsibility for interpretation. In a case discussing EPA's interpretation of one section of the Clean Air Act, the Court wrote:

> the regulatory scheme is technical and complex, the agency considered the matter in a detailed and reasoned fashion, and the decision involves reconciling conflicting policies.... Judges are not experts in the field, and are not part of either political branch of the Government.... When a challenge to an agency construction of a statutory provision, fairly conceptualized, really centers on the wisdom of the agency's policy, rather than whether it is a reasonable choice within a gap left open by Congress, the challenge must fail.[36]

Nevertheless, some scholars condemn the courts' default in favor of agency interpretation of statutes,[37,38] and thus in favor of agency policymaking authority. Other authors, however, insist that agencies are politically accountable and thus in the best position to make the policy decisions inherent in statutory interpretation.[39,40]

Of course, if the statute is fairly clear, a court will not allow an agency to disregard the statutory language, even though the statute vastly overregulates the risks. A good example of this occurred when a court reviewed an FDA decision to list as safe the color additives Orange No. 17 and Red No. 19, which are used in cosmetics. Although the agency determined, through animal bioassays, that these color additives were carcinogenic,[41] the calculated risk assessments showed that No. 17 would increase individual lifetime excess cancer risk by 2×10^{-10}, and No. 19 by 9×10^{-6}. The FDA concluded that these risks were too trivial to regulate. Indeed, there was some precedent to suggest that agencies possess an inherent statutory authority to disregard "de minimis" or trivial risks.[42] Unfortunately, this decision apparently conflicted with the express words of the Delaney Clause of the Food, Drug and Cosmetic Act, which prohibits all color additives that "induce cancer in man or animal." The court in *Public Citizen v. Young* recognized that the clause was "extraordinarily rigid," but it reluctantly adhered to the statutory wording.[43]

The other area of substantive judicial review involves the agency's judgment that a regulatory standard satisfies the "significant risk," "ample margin of safety," or some other statutory criterion. Under many statutes, a court may remand an agency decision only if it is "arbitrary, capricious or an abuse of discretion." Other statutes require the agency to show that its decision is supported by "substantial evidence." In interpreting these somewhat ambiguous standards the federal courts often have melded them into a single "hard-look" doctrine, which is a collection of techniques to control agency discretion. Under the doctrine, an agency decision can survive judicial review only if the agency has given a reasoned explanation of the bases for its decision, supported its decision with substantial evidence, explored alternatives, given reasons for rejecting the alternatives, and responded to public criticisms and objections.[44]

In determining whether the agency's decision is supported by substantial evidence, it is enough "that the administrative record contain(s) respectable scientific authority" supporting the agency's factual findings.[45] The Supreme Court has written:

> It is the Agency's responsibility to determine, in the first instance, what it considers to be a "significant" risk.... OSHA is not required to support its finding that a significant risk exists with anything approaching scientific certainty.... Thus, so long as they are supported by a body of reputable scientific thought, the Agency is free to use conservative assumptions in interpreting the data with respect to carcinogens, risking error on the side of overprotection rather than underprotection.[46]

Even where there are no data, or the data sharply conflict, a reviewing court will uphold the agency's policy judgments if the agency explains the considerations it relied on. In *Industrial Union Dep't, AFL-CIO v. Hodgson*, for example, the court readily upheld OSHA's decision to delay a new asbestos standard. Given the conflicting scientific views of the health impact of the delay, as well as the statutory policy to adopt feasible standards, the court held that the delay was not irrational.[47] Similarly, in *Building & Construction Trades Dep't, AFL-CIO v. Brock*, the court was unwilling to second-guess OSHA's finding that certain asbestos exposure levels posed a "significant risk" to workers.[48] In short, the courts generally defer to the agency's technical and administrative expertise even though the agency rationally might have made other, perhaps better, decisions.[49]

Inevitably, however, close judicial review of the agency's reasoning gives a court considerable opportunity to express its own substantive policy preferences. By immersing itself in the technical evidence, and determining whether the agency decision was "rational," the court often will not be able to avoid substituting its own views on the merits of the underlying issue. This may be especially true when the scientific evidence is most controverted. For example, in *Asbestos Information Ass'n v. OSHA*, the court reluctantly accepted the agency's assertion that an estimated eighty deaths, out of a worker population of 375,000, would constitute a "grave risk" that would justify an emergency temporary standard for asbestos exposure. Based on the court's own examination of the record, however, the court concluded that the numerical estimate was speculative and thus insufficient to support the agency's finding. As a result, the court stayed enforcement of the agency's emergency temporary standard.[50] Perhaps because it fears that lower courts will routinely intrude on agency policymaking authority, the Supreme Court in another case admonished that when reviewing agency decisions are "at the frontiers of science... a reviewing court must generally be at its most deferential."[51] Acceptability is more than probabilities and scientific assessments; it results from a decisionmaking process that is—on average—fair. When also viewed in the context of a portfolio of risks, the process tends to avoid ad hoc solutions.

The Portfolio Approach to Risk Acceptability

As illustrated in the preceding cases, risk assessment is based on a flux of scientific information concerning often highly uncertain or speculative risks. New knowledge, improved ability to control risks, and changes in risk attitudes can make a formerly acceptable risk no longer acceptable. Decisions about risk acceptability are thus dynamic and provisional; they must be monitored and adapted over time. For a regulatory agency, the acceptability of a specific risky activity depends on the context of other activities and control opportunities in which it is embedded. This view provides a new perspective for integrating risk acceptability issues into a framework for organizing social risk management decisions.[52]

An agency's approach to risk management can be viewed in terms of its management of four sets of risky activities. One set contains *known* problems waiting for regulation. A second set consists of *suspected* problems requiring further

investigation and possibly action. Finally, there are two disposition sets: one of *solved* problems that have been investigated and for which regulatory solutions have been established that must now be monitored and enforced; and one for *nonproblems* that have been investigated and found not to be problems. The known and suspected problem sets are sorted in order of decreasing priority of the problems in them, forming two rank-ordered lists. As suspected problems are investigated, and uncertainties about their risks are resolved, they change positions in the list for suspected problems. They may be moved down and off the suspected problem list altogether and onto the nonproblem list, or they may be moved up until they are pushed off the suspected problem list and inserted into the known problem list.

Risk management can be viewed as allocating resources to resolving problems on the two different action lists. In each budget period, the agency must allocate its resources to problems on the known and suspected lists. This requires trade-offs. Is it better to spend resources addressing another known problem or investigating another suspected one? When do the unknown risks from failing to explore items on the suspected problems list outweigh the losses from deferring action on known problems? And how should the pending problems within each list be ordered in terms of priority?

The only permanently acceptable risks are those moved to the nonproblem set. However, within any given period of time, there are known and suspected problems that are so far down on their corresponding priority sets that they will not be addressed until long after the many higher-priority problems that dominate them have been resolved. Such problems pose *provisionally acceptable* risks: risks that are acceptable until more important ones have been resolved.

In practice, new problems for investigation and resolution are continually being created by industrial society as new products and technologies emerge. Each new problem or potential problem requires positioning it, in the appropriate priority position, on the known problem or suspected problem list. If the rate at which high-priority problems are generated is greater than the rate at which they can be investigated, then existing problems will remain provisionally acceptable.

Of course, the process of setting the regulatory agenda is not as simple as this brief description suggests. Many problems are put on the public agenda as a result of the political and legal efforts of opposing parties. Rather than agency pull, in which problems are actively sought out for investigation to protect the public, public push may bring risk management problems onto the judicial, administrative, or legislative agenda. Episodic crises, such as the Bhopal disaster, also strongly affect the regulatory priorities regardless of the agency's assessment of their significance.

Although, in principle, toxic tort litigation can be included in the portfolio approach, the jurisdictional peculiarities of tort disputes stemming from state sovereignty, as well as the decentralized trial court system, prevent a centralized management of risks. This is less of a problem in the federal judicial system. Nevertheless, even though federal multidistrict litigation is possible, as the *Agent Orange* case demonstrates, tort law has a different focus than agency rulemaking under a risk management statute. Tort law emphasizes private, individual and class protection. Its objectives are principally compensation and to a lesser extent deterrence and occasionally punishment. Environmental statutes and agency rulemaking, by contrast, seek to prevent public injuries by enforcing standards adopted through a mixed political and technocratic process that may emphasize administrability as much as public health protection. Given these different objectives, it would be odd if the portfolio approach were equally applicable to both risk management systems.

The portfolio approach may also be inapplicable to toxic tort litigation because the individual risks, costs, and potential benefits of a case may differ from the societal ones. For example, since many of the costs of toxic tort disputes are borne by the affected parties, individual perceptions of risks, benefits and costs will affect the selection of disputes for litigation. The individual benefits of winning a case may be smaller than, or different from, the social benefits so that litigation that might be socially worthwhile may not be undertaken.[53] As a result, the portfolio approach is most difficult to apply to private litigation. Thus, social risk management by private litigation is a complement to, rather than a substitute for, risk management by regulatory agencies.

The concept of a portfolio of risks is also useful for companies whose allocation of resources—for example, to risk research, warning, and control—can be adapted to a common set of guidelines. Companies could use publicly stated levels of acceptable risk in their own risk management decisions. Without such clear guidance, manufacturers may be unwilling to produce socially beneficial but potentially risky products (e.g., vaccines) for fear of legal liability, should the courts decide in retrospect that the product was "unacceptably" risky.[54] An explicit reliance on the portfolio approach avoids the "carcinogen of the month" problem because it leads the agency expressly to consider items to be placed on the lists without becoming either engulfed by them or too lax. Public scrutiny during rulemaking and judicial review under regulatory law will now work toward plausible solutions to difficult and polycentric problems.

Conclusions

The concept of "acceptable" risk levels seems easiest to justify as a device for constraining and guiding regulatory risk management efforts and resource allocation over time. It is less clearly applicable to the private decisions of economic agents (consumers or producers), where the availability of detailed case-specific information about costs, uncertainties, and benefits makes it reasonable to expect and require more careful and detailed approaches to case-by-case risk management.

Although the need for simple, concrete, easily implemented standards of acceptable risks for guiding private sector health and safety risk management decisions cannot be denied, it seems unlikely that this need can be met without ignoring some important aspects of risk and uncertainty. The acceptability of risks is not an easy question and generally may not have answers that are both easy to apply and fully defensible on rational or moral grounds.

In this paper, we have argued that the most realistic and constructive view of risk acceptability for the practitioner may be that it is a property of risk management *decision processes*, rather than of isolated risky activities or situations. Numerous cases in the recent history of risk litigation indicate that it is the interaction of such processes—as in the balancing of the proper role of judicial intervention against the expertise and discretion of regulatory agencies—that determines risk acceptability in particular cases. Acceptability of a technological risk is thus not only a matter of risk statistics and objective numbers, but of social processes and of trade-offs that society is willing to make to achieve decisions that are *on average* reasonably fair, efficient, workable, and acceptable.

Acknowledgment

We gratefully acknowledge the stimulus and insightful comments of Professor Alvin W. Drake of M.I.T. The first two sections of this paper are an update of material in References 5 and 11. Paolo F. Ricci was in part supported by the Williams and Flora Hewlett Foundation to the Environmental Science and Engineering Program, UCLA.

References

1. P. F. Ricci, L. A. Cox, "Acceptability of chronic health risks," *Toxic Law Reporter* 986 (1987).
2. P. F. Ricci, L. S. Molton, "Health risk assessment: science, economics and law," *Annual Review of Energy* **11:** 77 (1986).
3. C. C. Travis, H. A. Hattemer-Frey, "Determining an acceptable level of risk," *Environ. Sci. Technol.* **22:** 873 (1988).
4. J. P. Dwyer, P. F. Ricci, "Coming to terms with acceptable risk," *Environ. Sci. Technol.* **23:** 145 (1989).
5. L. A. Cox, P. F. Ricci, "Risk, uncertainty and causation: quantifying human health risks," in *The Risk Assessment of Environmental Hazards*, D. J. Paustenbach, Ed., Wiley, New York, 1989, Chapter 2.
6. R. C. Elandt-Johnson, N. L. Johnson, *Survival Models and Data Analysis*, Wiley, New York, 1980.
7. S. D. Walter, "The estimation and interpretation of attributable risk in health research," *Biometrics* **32:** 829 (1976).
8. L. A. Cox, "Statistical issues in the estimation of assigned shares for carcinogenesis liability," *Risk Analysis* **7:** 81 (1987).
9. F. A. Seiler, H. K. Scott, "Mixtures of toxic agents and attributable risk calculations," *Risk Analysis* **7:** 61 (1987).
10. L. A. Cox, "Probability of causation and the attributable proportion of risk," *Risk Analysis* **4:** 221 (1984).
11. L. A. Cox, "Comparative risk measures for heterogeneous populations," in *Phenotypic Variation in Populations*, A. Woodhead et al., eds., Plenum, New York, 1988.
12. D. Maclean, "Social values and the distribution of risk," in *Values at Risk*, D. Maclean, ed., Rowman and Allanheld, Totowa, New Jersey, 1986.
13. P. C. Fishburn, P. D. Straffin, "Equity considerations in public risks evaluation," *Operations Research* **37:** 229 (1989).
14. B. Fischoff, S. Lichtenstein, R. Keeney, S. Derby, *Acceptable Risk*, Cambridge Univ. Press, London and New York, 1981.
15. R. L. Keeney, *Siting Energy Facilities*, Academic Press, New York, 1980.
16. L. A. Cox, "Economic theory of compensation rule design for probabilistic injuries," in L. B. Lave, Ed., *Risk Assessment and Management*, Plenum Press, New York, 1987.
17. *William Aldred's Case*, 77 Eng. Rep. 816 (1611).
18. J. Trauberman, "Statutory reform of 'toxic torts': relieving legal, scientific, and economic burdens on the chemical victim," *Harvard Environmental Law Review* **7:** 177 (1983).
19. O. F. Harris, "Toxic tort litigation and the causation element: is there any hope of reconciliation?" *Southwestern Law J.* **40:** 909 (1986).
20. R. L. Rabin, "Environmental liability and the tort system," *Houston Law Review* **24:** 27 (1987).
21. P. H. Schuck, "The new ideology of tort law," *The Public Interest* 93 (Summer 1988).
22. Comment, "The inapplicability of traditional tort analysis to environmental risks: the example of toxic waste pollution victim compensation," *Stanford Law Review* **35:** 575 (1983).
23. D. Rosenberg, "The causal connection in mass exposure cases: a 'public law' vision of the tort system," *Harvard Law Review* **97:** 849 (1984).
24. S. Gold, "Causation in toxic torts: burdens of proof, standards of persuasion, and statistical evidence," *Yale Law J.* **96:** 376 (1986).
25. T. A. Brennan, "Causal chains and statistical links: the role of scientific uncertainty in hazardous-substance litigation," *Cornell Law Review* **73:** 469 (1988).
26. *Sindell v. Abbott Laboratories*, 26 Cal. 3d 588, *cert. denied*, 449 U.S. 912 (1980).
27. P. H. Schuck, *Agent Orange on Trial: Mass Toxic Disasters in the Courts*, Harvard University Press, Cambridge, 1986.
28. *In re "Agent Orange" Product Liability Litigation*, 597 F. Supp. 740 (E.D.N.Y. 1984).
29. J. P. Dwyer, "Contentiousness and cooperation in environmental regulation," *American Journal of Comparative Law* **35:** 809 (1987).
30. C. S. Diver, "Policymaking paradigms in administrative law," *Harvard Law Review* **95:** 393 (1981).
31. *NRDC v. EPA*, 824 F.2d 1146 (D.C. Cir. 1987).
32. 53 Fed. Reg. 28,496 (1988); 54 Fed. Reg. 9,914 (1989).
33. N. A. Ashford, C. W. Ryan, C. C. Caldart, "Law and science policy in federal regulation of formaldehyde," *Science* **222:** 894 (1983).
34. *Industrial Union Dep't, AFL-CIO v. American Petroleum Institute*, 448 U.S. 607, 615 (1980).
35. *Id.* at 628, 645.
36. *Chevron v. NRDC*, 467 U.S. 837, 865-66 (1984).
37. J. R. Macey, "Promoting public-regarding legislation through statutory interpretation: an interest group model," *Columbia Law Review* **86:** 223 (1986).
38. W. N. Eskridge, "Politics without romance: implications of public choice theory for statutory interpretation," *Virginia Law Review* **74:** 275 (1988).
39. J. L. Mashaw, "Prodelegation: why administrators should make political decisions," *Journal of Law, Economics and Organization* **1:** 81 (1985).
40. C. S. Diver, "Statutory interpretation in the administrative state," *University of Pennsylvania Law Review* **133:** 549 (1985).
41. 51 Fed. Reg. 28,331 (1986); 51 Fed. Reg. 28,346 (1986).
42. *Alabama Power Co. v. Costle*, 636 F.2d 323 (D.C. Cir. 1979).
43. 831 F.2d 1108 (D.C. Cir. 1987).
44. *Ethyl Corp. v. EPA*, 541 F.2d 1, 34-36 (D.C. Cir.), *cert. denied*, 426 U.S. 941 (1976).
45. *EDF v. EPA*, 465 F.2d 528, 537 (D.C. Cir. 1972).
46. 448 U.S. at 655-56.
47. 499 F.2d 467, 474-79 (D.C. Cir. 1974).
48. 838 F.2d 1258, 1264-67 (D.C. Cir. 1988).
49. *Lead Industries Ass'n, Inc. v. EPA*, 647 F.2d 1130 (D.C. Cir.), *cert. denied*, 449 U.S. 1042 (1980).
50. 727 F.2d 415 (5th Cir. 1984).
51. *Baltimore Gas & Electric Co. v. NRDC*, 462 U.S. 87, 103 (1983).
52. L. A. Cox, P. F. Ricci, "Legal and philosophical aspects of risk analysis," in *The Risk Assessment of Environmental Hazards*, D. J. Paustenbach, ed., Wiley, New York, 1989, Chapter 30.
53. S. Shavell, "A model of the optimal use of liability and safety regulation," *Rand Journal of Economics* **15:** 271 (1984).
54. K. S. Abraham, R. A. Merrill, "Scientific uncertainty in the courts," *Issues in Science and Technology* 93 (Winter 1986).

Paolo F. Ricci is an Associate Professor at UCLA, School of Public Health, 46-078 CHS, Los Angeles, CA 90024. He is also an Adj. Professor of Law at U.C. Berkeley, CA. Louis Anthony Cox, Jr. is affiliated with U.S. West Advanced Technologies, Englewood, Colorado. John P. Dwyer is an Acting Professor of Law, at the University of California, Berkeley, California.

This paper was submitted for peer review on May 5, 1988. The revised manuscript was received March 20, 1989.